arXiv:1704.06299v1 [math.CO] 20 Apr 2017

Complexity of the Fourier transform on the Johnson graph

Rodrigo Iglesias
Departamento de Matemática
Universidad Nacional del Sur
Bahía Blanca, Argentina.

Mauro Natale
Facultad de Ciencias Exactas
Universidad Nacional del Centro
Tandil, Argentina.

April 24, 2017

Abstract

The set X of k-subsets of an n-set has a natural graph structure where two k-subsets are connected if and only if the size of their intersection is $k-1$. This is known as the Johnson graph. The symmetric group S_n acts on the space of complex functions on X and this space has a multiplicity-free decomposition as sum of irreducible representations of S_n, so it has a well-defined Gelfand-Tsetlin basis up to scalars. The Fourier transform on the Johnson graph is defined as the change of basis matrix from the delta function basis to the Gelfand-Tsetlin basis.

The direct application of this matrix to a generic vector requires $\binom{n}{k}^2$ arithmetic operations. We show that –in analogy with the standard Fast Fourier Transform on the discrete circle– this matrix can be factorized as a product of $n-1$ orthogonal matrices, each one with at most two nonzero elements in each column. This factorization shows that the number of arithmetic operations required to apply this matrix to a generic vector is bounded above by $2(n-1)\binom{n}{k}$.

The proof is based on a restricted version of the Robinson-Schensted algorithm. In effect, it is based on the construction of $n-1$ intermediate bases, each one parametrized by certain pairs composed by a standard Young tableau and a word.

1 Introduction

The set of all subsets of cardinal k of a set of cardinal n is a basic combinatorial object with a natural metric space structure where two k-subsets are at distance d if the size of their intersection is $k-d$. This structure is captured by the Johnson graph, whose nodes are the k-subsets and two k-subsets are connected if and only if they are at distance 1.

The Johnson graph is closely related to the Johnson scheme, an association scheme. It has been also involved in the graph isomorphism problem. It emerged as a hard case that caused numerous graph isomorphism algorithms to fail. Recently, the Johnson graph played a fundamental role in the breakthrough quasipolynomial time algorithm for the graph isomorphism problem [4].

The metric structure of the Johnson graph raises the intuitive notion of smooth function on the set of k-sets: a function is smooth if every two nearby k-sets have similar

values. In the same way that a smooth function on the discrete circle can be well approximated by a sum of a few low frequency exponentials via the Discrete Fourier Transform (DFT), a smooth function on the set of k-sets is expected to be well approximated by a sum of a few special functions via a suitable non-abelian Fourier transform. This suggests that many interesting functions and distributions on the set of k-sets can be compressed by representing them as linear combination of these special functions. This motivates the question of how this special functions are defined and how the decomposition of a given function as sum of the special functions can be computed efficiently.

Functions on the Johnson graph arise in the analysis of ranked data. In many contexts, a big number of agents choose a k-subset from an n-set, and the data is collected as the function that assigns to the k-subset x the number of agents who choose x. This situation is considered, for example, in the statistical analysis of certain lotteries (see [3]).

The vector space of functions on the Johnson graph is a representation of the symmetric group and it decomposes as a multiplicity-free direct sum of irreducible representations. Statistically relevant information about the function is contained in the isotypic projections of the function onto each irreducible component. This approach to the analysis of ranked data was called spectral analysis by Diaconis and developed in [2], [3]. The problem of the efficient computation of the isotypic projections has been studied by Diaconis and Rockmore in [5], and by Maslen, Orrison and Rockmore in [8].

The classical DFT on the cyclic group $\mathbb{Z}_N$ can be seen as the application of a change of basis matrix from the basis B_0 of delta functions to the basis B_N of characters of the group $\mathbb{Z}_N$. The direct application of this matrix to a generic vector involves N^2 arithmetic operations. The Fast Fourier Transform (FFT) is a fundamental algorithm dicovered by Cooley and Tukey [1] that computes the DFT in $O(N \log N)$ operations. This efficiency is due to a factorization of the change of basis matrix

$$[B_0]_{B_N} = [B_{N-1}]_{B_N} \ ... \ [B_1]_{B_2} \ [B_0]_{B_1}$$

where $B_1, ..., B_{N-1}$ are intermediate orthonormal bases such that each matrix $[B_{i-1}]_{B_i}$ has two nonzero entries in each column.

In this paper, we show that the same phenomenon occurs in the case of the nonabelian Fourier transform on the Johnson graph. The argument is based on the structure of the set of standard tableaux of height at most two and a restricted version of the Robinson-Schensted correspondence.

The Fourier transform on the Johnson graph is defined as the application of the change of basis matrix from the basis B_0 of delta functions to the basis B_n of Gelfand-Tsetlin functions. The Gelfand-Tsetlin basis –defined in Section 2– is well-behaved with respect to the action of the symmetric group S_n, in the sense that each irreducible component is generated by a subset of the basis.

A direct computation of this Fourier transform involves $\binom{n}{k}^2$ arithmetic operations. The computational model used here counts a single complex multiplication and addition as one operation. We construct intermediate orthonormal bases $B_1, ..., B_{n-1}$ such that each change of basis matrix $[B_{i-1}]_{B_i}$ has at most two nonzero entries in each column. Each intermediate basis B_i is parametrized by pairs composed by a standard Young tableau of height at most two and a word in the alphabet $\{1, 2\}$ as shown in Figure 1.

These intermediate bases enable to compute the nonabelian Fourier transform –as well as its inverse– in at most $2(n-1)\binom{n}{k}$ operations.

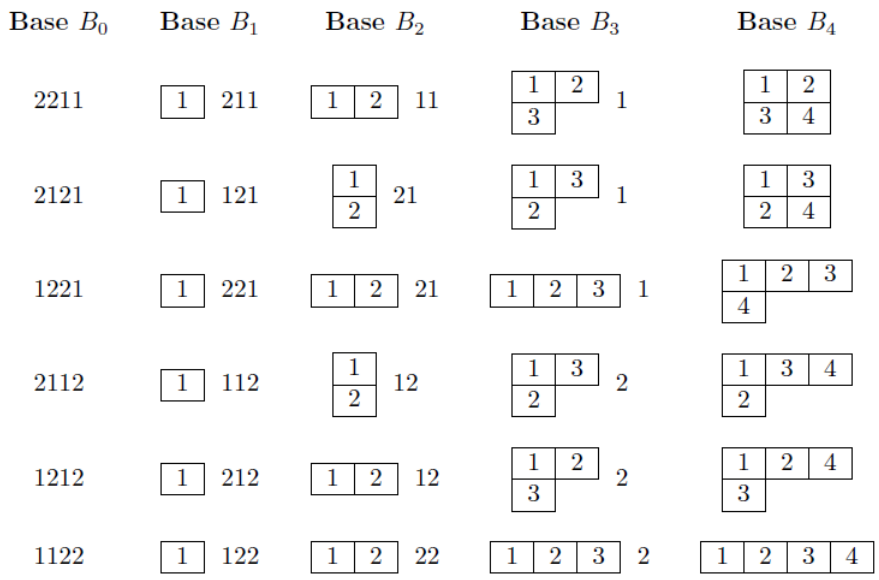

Figure 1: Labels of the intermediate bases in the case $n = 4, k = 2$. The i-th column parametrize the basis B_i. Each row shows a process that is a consequence of the Robinson-Schensted algorithm.

In Section 2, we review the definition of Gelfand-Tsetlin bases for representations of the symmetric group. In Section 3, we describe the well-known decomposition of the function space on the Johnson graph and define the corresponding Gelfand-Tsetlin basis. In Section 4, we introduce the sequence of intermediate bases of the function space, analyze the sparsity of the change of basis matrix between two consecutive bases and prove the main result, contained in Theorem 2. In Section 5, we mention some applications of our upper bound to the problem of the computation of the isotypic components of a function.

2 Gelfand-Tsetlin bases

Consider the chain of subgroups of S_n

$$\{Id\} \subset S_2 \subset S_3 ... \subset S_n$$

where S_k is the subgroup of those permutations fixing the last $n-k$ elements of $\{1, ..., n\}$. Let $Irr(n)$ be the set of equivalence classes of irreducible representations of S_n. A fundamental fact in the representation theory of S_n is that if V_λ is an irreducible S_n-module corresponding to the representation $\lambda \in Irr(n)$ and we consider it by restriction as an S_{n-1}-module, then it decomposes as sum of irreducible representation of S_{n-1} in a multiplicity-free way (see for example [10]). This means that if V_μ is an irreducible S_{n-1}-module corresponding to the representation $\mu \in Irr(n-1)$ then

$$\dim \mathrm{Hom}_{S_{n-1}}(V_\mu, V_\lambda)$$

is 0 or 1. The *branching graph* is the following directed graph. The set of nodes is the disjoint union

$$\bigsqcup_{n\geq 0} Irr(n)$$

Given representations $\lambda \in Irr(n)$ and $\mu \in Irr(n-1)$ there is an edge connecting them if and only if μ appears in the decomposition of λ, that is, if $\dim \mathrm{Hom}_{S_{n-1}}(V_\mu, V_\lambda) = 1$. If there is an edge between them we write

$$\mu \nearrow \lambda$$

so we have a canonical decomposition into irreducible S_{n-1}-modules

$$V_\lambda = \bigoplus_{\mu \nearrow \lambda} V_\mu$$

Applying this formula iteratively we obtain a uniquely determined decomposition into one-dimensional subspaces

$$V_\lambda = \bigoplus_{T} V_T$$

where T runs over all chains

$$T = \lambda_1 \nearrow \lambda_2 \nearrow ... \nearrow \lambda_n$$

where $\lambda_i \in Irr(i)$ and $\lambda_n = \lambda$. Choosing a unit vector v_T of the one-dimensional space V_T we obtain a basis $\{v_T\}$ of the irreducible module V_λ, which is called the *Gelfand-Tsetlin basis*.

Observe that if V is a multiplicity-free representation of S_n then there is a uniquely determined –up to scalars– Gelfand-Tsetlin basis of V. In effect, if

$$V = \bigoplus_{\lambda \in S \subseteq Irr(n)} V_\lambda$$

and B_λ is a GT-basis of V_λ then a GT-basis of V is given by the disjoint union

$$B = \bigsqcup_{\lambda \in S \subseteq Irr(n)} B_\lambda$$

The Young graph is the graph where the nodes are the Young diagrams and there is an arrow from λ to μ if and only if λ is contained in μ and their difference consists in only one box. It turns out that the branching graph is isomorphic to the Young graph and there is a bijection between the set of Young diagrams with n boxes and $Irr(n)$.

Then there is a bijection between the Gelfand-Tsetlin basis of V_λ (where λ is a Young diagram) and the set of paths in the Young graph starting at the one-box diagram and ending at the diagram λ. Each path can be represented by a unique standard Young tableau, so that the Gelfand-Tsetlin-basis of V_λ is parametrized by the set of standard Young tableaux of shape λ. From now on we identify a chain $\lambda_1 \nearrow \lambda_2 \nearrow ... \nearrow \lambda_n$ with its corresponding standard tableau.

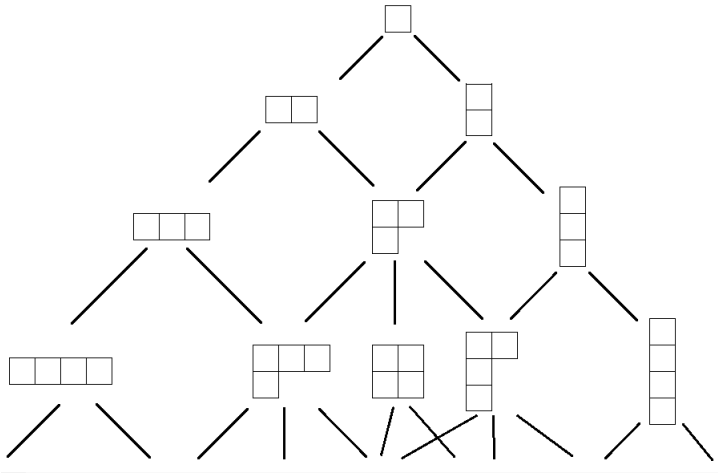

Figure 2: The Young graph. Each path from the top node to a particular Young diagram is identified with a standard Young tableau.

3 Decomposition of the function space on the Johnson graph

A k-set is a subset of $\{1, ..., n\}$ of size k and let X be the set of all k-sets. Given two k-sets x, y the distance $d(x, y)$ is defined as $n - |x \cap y|$. The group S_n acts naturally on X by

$$\sigma\{i_1, ..., i_k\} = \{\sigma(i_1), ..., \sigma(i_k)\}$$

The vector space $\mathcal{F}$ of the complex valued functions on X is a representation of S_n where the action is given by $\sigma f = f \circ \sigma^{-1}$. The representation $\mathcal{F}$ decomposes uniquely as a direct sum

$$\mathcal{F} = \bigoplus_{\lambda \in Irr(n)} n_\lambda \, V_\lambda$$

where n_λ is the number of copies of the irreducible representation V_λ. Given a function $f \in \mathcal{F}$ there is a unique decomposition

$$f = \sum_\lambda f_\lambda$$

with $f_\lambda \in n_\lambda V_\lambda$. We say that f_λ is the projection of f onto the isotypic component $n_\lambda V_\lambda$ and let $P_\lambda : \mathcal{F} \to \mathcal{F}$ be this orthogonal projection. These projections form a resolution of the identity

$$I = \sum_\lambda P_\lambda, \qquad\qquad P_\lambda P_\mu = \delta_{\lambda\mu} P_\lambda.$$

There is a general formula to compute the projections P_λ. Let G be a finite group acting on a vector space V. A linear endomorphism $T : V \to V$ is called a G-map if it is G-invariant, that is, $gT = Tg$ for all $g \in G$. A function $\phi : G \to \mathbb{C}$ is called a class function if it satisfies $\phi(hgh^{-1}) = \phi(g)$ for all $h, g \in G$.

Proposition 1. *Let V be a representation of G and let $\phi : G \to \mathbb{C}$ be a class function. Then the space $End(V)$ of all endomorphisms of the form*

$$\frac{1}{|G|}\sum_{g\in G}\phi(g)\rho_V(g)$$

is a commutative subalgebra of the algebra of G-maps. The isotypic projections form a linear base $\{P_\lambda\}_{\lambda\in Irr(G)}$ of this subalgebra and they are obtained by setting $\phi = \bar{\chi_\lambda}$, where χ_λ is the character of the irreducible representation λ.

For a proof see, for example, [7]. In our situation, V is the space $\mathcal{F}$ of functions on the Johnson graph so V has dimension $\binom{n}{k}$ and the group $G = S_n$ has cardinal $n!$. In this case, the size of the group G is exponential in the dimension of V so the formula is unfeasible. This suggests that to obtain more practical formulas for the projections we need to explore simple G-maps requiring only a small number of arithmetical operations and try to express the projections P_λ in terms of them.

In order to obtain simple G-maps, we search for endomorphisms having a small orbit. Let $\{b_s\}_{s\in X}$ be the delta basis of $V = \mathcal{F}$ and let $\{E_{st}\}_{s,t\in X}$ be the base of $End(\mathcal{F})$ where E_{st} is the endomorphism sending b_s to b_t and sending the rest of the base to 0. Observe that the orbit of $\{E_{st}\}$ consists of all the elements $\{E_{s't'}\}$ such that $d(s',t') = d(s,t)$. Symmetrizing these elements we obtain

$$D_r = \sum_{dist(s,t)=r} E_{st}$$

Then the set $\{D_r\}_{r=0,\dots,k}$ forms an orthogonal basis of the space $End(\mathcal{F})^G$ of all the G-maps, where $G = S_n$. In particular, we see that the dimension of $End(\mathcal{F})^G$ is $k+1$.

One of the simplest elements of this base is D_1 because it has a small orbit. Each k-set has $k(n-k)$ neighbours at distance 1. Then the matrix of D_1 in the canonical base is a symmetric 0-1 matrix where each row has $k(n-k)$ nonzero elements. Applying this matrix to a vector requires at most $k(n-k)\,|X|$ arithmetic operations.

The operator D_r satisfies the recurrence

$$D_1 D_r = c_{r+1}D_{r+1} + a_r D_r + b_{r-1}D_{r-1}$$

where

$$\begin{aligned}
c_{r+1} &= |\{z \mid d(x,z)=1, d(z,y)=r\}| \text{ for any fixed } x,y \text{ with } d(x,y)=r+1.\\
a_r &= |\{z \mid d(x,z)=1, d(z,y)=r\}| \text{ for any fixed } x,y \text{ with } d(x,y)=r.\\
b_{r-1} &= |\{z \mid d(x,z)=1, d(z,y)=r\}| \text{ for any fixed } x,y \text{ with } d(x,y)=r-1.
\end{aligned}$$

As a consequence, the operator D_r is a polynomial in D_1

$$D_r = p_r(D_1)$$

In particular, this shows that the algebra of G-maps is commutative. It is a general fact that the algebra of G-maps is commutative if and only if the decomposition of V is multiplicity-free. This proves the following.

Proposition 2. *The space $\mathcal{F}$ of functions on the Johnson graph decomposes in $k+1$ multiplicity-free irreducible representations of the group S_n*

$$\mathcal{F} = \bigoplus_{i=0}^{k} V_{\lambda_i}$$

Then the space $\mathcal{F}$ has a well-defined –up to scalars– Gelfand-Tsetlin basis. We also observe that the operator P_{λ_i} that orthogonally projects onto the component V_{λ_i} is a G-map and the set $\{P_{\lambda_i}\}_{i=0}^{k}$ is a linear base of the algebra of G-maps.

3.1 Gelfand-Tsetlin basis of $\mathcal{F}$

The shape of a Young diagram is the sequence of numbers of boxes in the rows, from bottom to top. For example the shape of

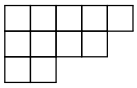

is $(2,4,5)$. It can be shown (see [9]) that the decomposition of $\mathcal{F}$ is given as follows.

Theorem 1. *The $k+1$ irreducible representations of S_n appearing in the decomposition of the space of functions on the Johnson graph are given by the Young diagrams of shape $(n),(1,n-1),(2,n-2),...,(k,n-k)$.*

For example, if $n=6$ and $k=2$ then

$\mathcal{F}$ = 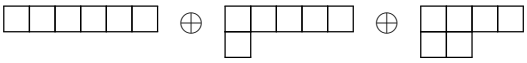

This result gives a parametrization of the GT-basis of $\mathcal{F}$: there is a bijection between the set of elements of the GT-basis and the set of standard tableaux of shape $(a,n-a)$ where a runs from 0 to k.

Let us give a more explicit description of the GT-basis of $\mathcal{F}$. Consider the space $\mathcal{F}$ as an S_i-module for $i=1,...,n$, and let $\mathcal{F}_{i,\lambda}$ be the isotypic component corresponding to the irreducible representation λ of S_i so that for each i we have a decomposition

$$\mathcal{F} = \bigoplus_{\lambda \in Irr(S_i)} \mathcal{F}_{i,\lambda}$$

where $\mathcal{F}_{i,\lambda} \perp \mathcal{F}_{i,\lambda'}$ if $\lambda \neq \lambda'$. For each standard tableau $\lambda_1 \nearrow \lambda_2 \nearrow ... \nearrow \lambda_n$ let

$$\mathcal{F}_{\lambda_1 \nearrow \lambda_2 \nearrow ... \nearrow \lambda_n} = \mathcal{F}_{1,\lambda_1} \cap \mathcal{F}_{2,\lambda_2} \cap ... \cap \mathcal{F}_{n,\lambda_n}$$

Then Theorem 1 shows that $\mathcal{F}$ has an orthogonal decomposition in one-dimensional subspaces

$$\mathcal{F} = \bigoplus_{\lambda_1 \nearrow \lambda_2 \nearrow ... \nearrow \lambda_n} \mathcal{F}_{\lambda_1 \nearrow \lambda_2 \nearrow ... \nearrow \lambda_n}$$

where λ_n runs through all representations of S_n corresponding to Young diagrams of shape $(a,n-a)$ for $a=0,...,k$.

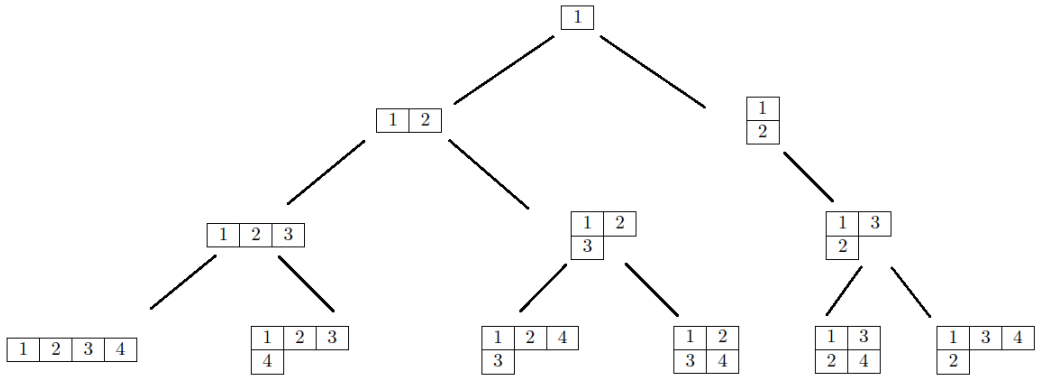

Figure 3: The leaves of this tree parametrize the Gelfand-Tsetlin basis of the space $\mathcal{F}$ of functions on the Johnson graph $J(4,2)$.

4 The intermediate bases

Let us describe schematically the Fast Fourier Transform algorithm for the Johnson graph. The input is a vector f in the space $\mathcal{F}$ of functions on the space of k-sets X, written in the delta function basis B_0, given as a column vector $[f]_{B_0}$. The output of the algorithm is a column vector representing the vector f written in the basis B_n, the Gelfand-Tsetlin basis of $\mathcal{F}$. In other words, the objective is to apply the change of basis matrix to a given column vector:

$$[f]_{B_n} = [B_0]_{B_n}\ [f]_{B_0}$$

Our technique to realize this matrix multiplication is to construct a sequence of intermediate orthonormal bases $B_1, B_2, ..., B_{n-1}$ such that

$$[B_0]_{B_n} = [B_0]_{B_1}\ \ [B_1]_{B_2}\ \ \ldots\ \ [B_{n-1}]_{B_n}$$

is a decomposition where each factor is a very sparse matrix.

4.1 Definition of the basis B_i

We represent a k-set by a word in the alphabet $\{1,2\}$ as follows. The element $i \in \{1,...,n\}$ belongs to the k-subset if and only if the place i of the word is occupied by the letter 1. For example,

$$\{2,3,6,8\} \subseteq \{1,...,9\} \qquad \to \qquad 2\ 1\ 1\ 2\ 2\ 1\ 2\ 1\ 2$$

So, from now on, we identify X with the set of words of length n in the alphabet $\{1,2\}$ such that the letter 1 appears k times. The group S_n acts on X in the natural way.

The subgroup S_i with $1 \leq i \leq n$ acts on the first i letters fixing the last $n-i$ letters of the word. A word w of length $n-i$ is called an *i-tail* if there is a word w' of length i such that the concatenation $w'w$ is a word in X. Let X_i be the set of i-tails.

For each i-tail $w \in X_i$ we define X^w as the subset of those words $x \in X$ with i-tail w, that is, those words x that consist of a concatenation $x = w'w$ for some word w' of length

i. Observe that each subset X^w is stabilized by the action of the subgroup S_i. To each word $x \in X$ we attach the delta function $b(x)$ defined on X by

$$b(x)(z) = \begin{cases} 1 & \text{if } x = z \\ 0 & \text{otherwise} \end{cases}$$

Let $\mathcal{F}^w$ be the subspace of $\mathcal{F}$ spanned by the functions $b(x)$ such that $x \in X^w$:

$$\mathcal{F}^w = \bigoplus_{x \in X^w} \mathbb{C}\ b(x)$$

Then $\mathcal{F}$ decomposes as

$$\mathcal{F} = \bigoplus_{w \in X_i} \mathcal{F}^w$$

and each subspace $\mathcal{F}^w$ is invariant by the action of S_i.

The following is a key observation. Suppose that the letter 1 appears r times in the tail w. Then the subset X^w consists of those words of the form $w'w$ such that w' is of length i and it has exactly $k - r$ appearances of the letter 1. This means that X^w has the structure of the Johnson graph $J(i, k-r)$ and, when acted by the subgroup S_i, the space of $\mathbb{C}$-valued functions on X^w decomposes as an S_i-module in a multiplicity-free way according to the formula of Theorem 1.

As a consequence, each subspace $\mathcal{F}^w$ has a Gelfand-Tsetlin decomposition

$$\mathcal{F}^w = \bigoplus_{\lambda_1 \nearrow \lambda_2 \nearrow ... \nearrow \lambda_i} \mathcal{F}^w_{\lambda_1 \nearrow \lambda_2 \nearrow ... \nearrow \lambda_i}$$

where λ_i runs over all Young diagrams of shape $(a, i-a)$ with $0 \leq a \leq k - r$. Then $\mathcal{F}^w$ has a Gelfand-Tsetlin basis B^w, uniquely determined up to scalars. We define, up to scalars, the i-th intermediate basis of $\mathcal{F}$ as

$$B_i = \bigsqcup_{w \in X_i} B^w$$

From Theorem 1, we see that the basis B^w is parametrized by the set of standard tableaux of shape $(a, i-a)$ with $0 \leq a \leq k-r$. On the other hand, the word w runs over the set of i-tails. Figure 4 illustrates the structure of the intermediate bases.

4.2 Sparsity of the change of basis matrix $[B_i]_{B_{i-1}}$

Let $\mathcal{F}_{i,\lambda}$ be defined as in Section 3.1. On the other hand, for $i = 1, ..., n$ and $k \in \{1, 2\}$ let us define $\mathcal{F}^{i,k}$ as the subspace of $\mathcal{F}$ generated by the delta functions $b(x)$ such that the word x has the letter k in the place i. For each i we have a decomposition

$$\mathcal{F} = \mathcal{F}^{i,1} \oplus \mathcal{F}^{i,2}$$

where $\mathcal{F}^{i,1} \perp \mathcal{F}^{i,2}$. Let w be an i-tail

$$w = k_{i+1} k_{i+2} ... k_n$$

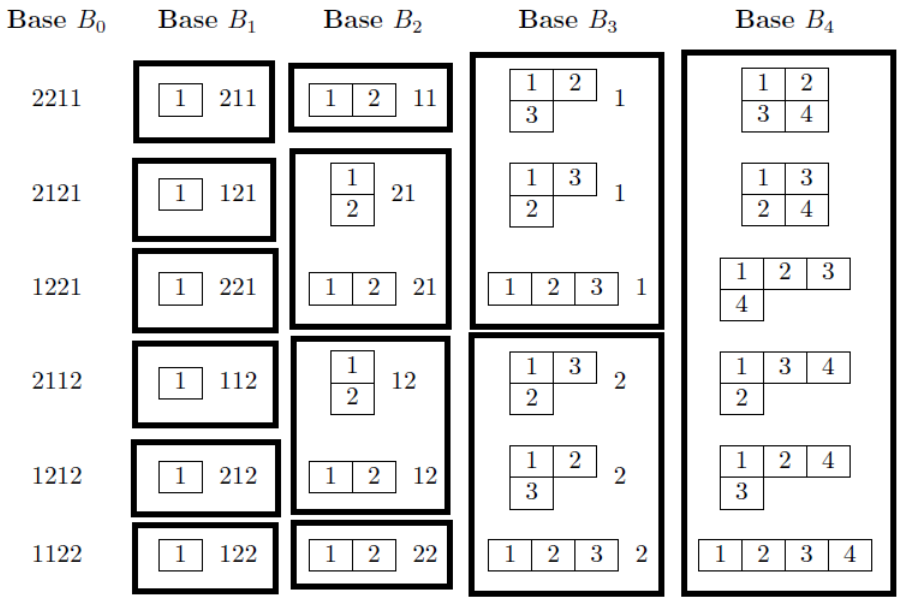

Figure 4: Labels of the intermediate bases in the case $n = 4, k = 2$. The boxes in each column represent the decomposition $B_i = \sqcup_{w \in X_i} B^w$.

Then

$$\mathcal{F}^w = \mathcal{F}^{i+1,k_{i+1}} \cap \ldots \cap \mathcal{F}^{n-1,k_{n-1}} \cap \mathcal{F}^{n,k_n}$$

and the one-dimensional subspaces generated by elements of the base B_i can be expressed as

$$\mathcal{F}^w_{\lambda_1 \nearrow \lambda_2 \nearrow \ldots \nearrow \lambda_i} = \mathcal{F}_{1,\lambda_1} \cap \mathcal{F}_{2,\lambda_2} \cap \ldots \cap \mathcal{F}_{i,\lambda_i} \cap \mathcal{F}^{i+1,k_{i+1}} \cap \ldots \cap \mathcal{F}^{n-1,k_{n-1}} \cap \mathcal{F}^{n,k_n}$$

Let us compare the base B_i with the base B_{i-1}.

Lemma 1. *Suppose that the subspace*

$$\mathcal{F}_{1,\lambda_1} \cap \ldots \cap \mathcal{F}_{i-1,\lambda_{i-1}} \cap \mathcal{F}_{i,\lambda_i} \cap \mathcal{F}^{i+1,k_{i+1}} \cap \ldots \cap \mathcal{F}^{n,k_n}$$

generated by an element of the basis B_i is not orthogonal to the subspace

$$\mathcal{F}_{1,\lambda'_1} \cap \ldots \cap \mathcal{F}_{i-1,\lambda'_{i-1}} \cap \mathcal{F}^{i,k'_i} \cap \mathcal{F}^{i+1,k'_{i+1}} \cap \ldots \cap \mathcal{F}^{n,k'_n}$$

generated by an element of the basis B_{i-1}. Then

$$\lambda_1 = \lambda'_1, \quad \ldots \quad , \lambda_{i-1} = \lambda'_{i-1} \qquad \textit{and} \qquad k_{i+1} = k'_{i+1}, \quad \ldots \quad , k_n = k'_n$$

Proof. If $\lambda_j \neq \lambda'_j$ for some j with $1 \leq j \leq i-1$, then

$$\mathcal{F}_{j,\lambda_j} \perp \mathcal{F}_{j,\lambda'_j}$$

and the two subspaces are orthogonal. If $k_j \neq k'_j$ for some j with $i+1 \leq j \leq n$, then

$$\mathcal{F}^{j,k_j} \perp \mathcal{F}^{j,k'_j}$$

and the two subspaces are orthogonal. □

Corollary 1. *For each element b of the base B_i there is at most two elements of the base B_{i-1} not orthogonal to b.*

Proof. Suppose that the element b of B_i generates the subspace

$$\mathcal{F}_{1,\lambda_1} \cap ... \cap \mathcal{F}_{i-1,\lambda_{i-1}} \cap \mathcal{F}_{i,\lambda_i} \cap \mathcal{F}^{i+1,k_{i+1}} \cap ... \cap \mathcal{F}^{n,k_n}$$

By Lemma 1, the subspace generated by an element of B_{i-1} not orthogonal to b is of the form

$$\mathcal{F}_{1,\lambda_1} \cap ... \cap \mathcal{F}_{i-1,\lambda_{i-1}} \cap \mathcal{F}^{i,k_i} \cap \mathcal{F}^{i+1,k_{i+1}} \cap ... \cap \mathcal{F}^{n,k_n}$$

where $k_i = 1$ or $k_i = 2$. □

Theorem 2. *Let B_0 be the delta function basis of $\mathcal{F}$ and let B_n be a Gelfand-Tsetlin basis of $\mathcal{F}$. Given a column vector $[f]_{B_0}$ with $f \in \mathcal{F}$, the column vector $[f]_{B_n}$ given by*

$$[f]_{B_n} = [B_0]_{B_n}\ [f]_{B_0}$$

can be computed using no more than $2(n-1)\binom{n}{k}$ operations.

Proof. By Corollary 1 we see that each column of the matrix $[B_{i-1}]_{B_i}$ has at most two non-zero elements. Then the application of the matrix $[B_{i-1}]_{B_i}$ to a generic column vector can be done using at most $2\binom{n}{k}$ operations. Observe that $[B_0]_{B_1}$ is the identity matrix. We have

$$\begin{aligned}[B_0]_{B_n} &= [B_0]_{B_1}\ [B_1]_{B_2}\ ...\ [B_{n-1}]_{B_n} \\ &= [B_1]_{B_2}\ ...\ [B_{n-1}]_{B_n}\end{aligned}$$

Then the successive applications of the $n-1$ matrices can be done in at most $2(n-1)\binom{n}{k}$. □

4.3 Example

In the case $n = 4$, $k = 2$, the matrices $[B_i]_{B_{i-1}}$ in the decomposition of $[B_0]_{B_4}$ are the following:

$$[B_1]_{B_2} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & -\frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} & 0 & 0 & 0 \\ 0 & \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} & 0 & 0 & 0 \\ 0 & 0 & 0 & -\frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} & 0 \\ 0 & 0 & 0 & \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

$$[B_2]_{B_3} = \begin{bmatrix} -\frac{\sqrt{6}}{3} & 0 & \frac{\sqrt{3}}{3} & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ \frac{\sqrt{3}}{3} & 0 & \frac{\sqrt{6}}{3} & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & -\frac{\sqrt{3}}{3} & \frac{\sqrt{6}}{3} \\ 0 & 0 & 0 & 0 & \frac{\sqrt{6}}{3} & \frac{\sqrt{3}}{3} \end{bmatrix}$$

$$[B_3]_{B_4} = \begin{bmatrix} -\frac{\sqrt{2}}{2} & 0 & 0 & 0 & \frac{\sqrt{2}}{2} & 0 \\ 0 & -\frac{\sqrt{2}}{2} & 0 & \frac{\sqrt{2}}{2} & 0 & 0 \\ 0 & 0 & -\frac{\sqrt{2}}{2} & 0 & 0 & \frac{\sqrt{2}}{2} \\ 0 & \frac{\sqrt{2}}{2} & 0 & \frac{\sqrt{2}}{2} & 0 & 0 \\ \frac{\sqrt{2}}{2} & 0 & 0 & 0 & \frac{\sqrt{2}}{2} & 0 \\ 0 & 0 & \frac{\sqrt{2}}{2} & 0 & 0 & \frac{\sqrt{2}}{2} \end{bmatrix}$$

5 Applications

The upper bound we obtained for the algebraic complexity of the Fourier transform on the Johnson graph can be applied to the problem of computing all the weights of the irreducible components of a given function. In effect –after the precomputation of the n intermediate matrices– all the lengths of the isotypic projections of the original column vector can be directly computed from the tranformed vector, so the problem can be solved in $O(n\binom{n}{k})$ operations. This problem could also be solved by computing every isotypic component and measuring their lengths, but the most efficient algorithm to compute all the isotypic projections –given by Maslen, Orrison and Rockmore in [8]– uses $O(k^2\binom{n}{k})$ operations, so our upper bound is an improvement when k asymptotically dominates $\sqrt{n}$.

Similarly, we can apply the upper bound to the problem of computing the isotypic projection onto a particular set of components. Actually, we can perform the projection as follows: first we apply the Fourier transform, then we annihilate the coordinates corresponding to the irreducible components not included in the set, and finally we apply the inverse Fourier transform. Then the projection can be performed in at most $4(n-1)\binom{n}{k}$ operations.

References

[1] Cooley, James W. and Tukey, John W. "An Algorithm for the Machine Calculation of Complex Fourier Series." IBM Watson Research Center (1964): 297-301.

[2] Diaconis, Persi. "A generalization of spectral analysis with application to ranked data." The Annals of Statistics (1989): 949-979.

[3] Diaconis, Persi. "Group representations in probability and statistics." Lecture Notes-Monograph Series (1988): i-192.

[4] Babai, Laszlo. "Graph Isomorphism in Quasipolynomial Time." arXiv:1512.03547v1 (2015)

[5] Diaconis, Persi, and Daniel Rockmore. "Efficient computation of isotypic projections for the symmetric group." DIMACS Ser. Discrete Math. Theoret. Comput. Sci 11 (1993): 87-104.

[6] Driscoll, James R., Dennis M. Healy Jr, and Daniel N. Rockmore. "Fast discrete polynomial transforms with applications to data analysis for distance transitive graphs." SIAM Journal on Computing 26.4 (1997): 1066-1099.

[7] Fulton, William, and Joe Harris. Representation theory. Vol. 129. Springer Science and Business Media, 1991.

[8] Maslen, David K., Michael E. Orrison, and Daniel N. Rockmore. "Computing isotypic projections with the Lanczos iteration." SIAM Journal on Matrix Analysis and Applications 25.3 (2003): 784-803.

[9] Stanton, Dennis. "Orthogonal polynomials and Chevalley groups." Special Functions: group theoretical aspects and applications. Springer Netherlands, 1984. 87-128.

[10] Vershik, Andrei. M. and Okounkov Anatoly. Yu. "A New Approach to Representation Theory of the Symmetric Groups." ESI The Erwin Schödinger International Institute for Mathematical Physics (1996): 1-21.